

# (12) UK Patent Application (19) GB (11) 2 276 962 (13) A

(43) Date of A Publication 12.10.1994

(21) Application No 9404866.7

(22) Date of Filing 14.03.1994

(30) Priority Data

(31) 9307359

(32) 08.04.1993

(33) GB

(71) Applicant(s)

**International Computers Limited**

**(Incorporated in the United Kingdom)**

**ICL House, 1 High Street, Putney, LONDON,  
SW15 1SW, United Kingdom**

(72) Inventor(s)

**Albert Stephen Hilditch**

(74) Agent and/or Address for Service

**D C Guyatt**

**International Computers Limited, Intellectual  
Property Dept, Cavendish Road, STEVENAGE,  
Hertfordshire, SG1 2DY, United Kingdom**

(51) INT CL<sup>5</sup>

**G06F 12/12**

(52) UK CL (Edition M )

**G4A AMC**

(56) Documents Cited

**EP 0391871 A2**

(58) Field of Search

UK CL (Edition M ) **G4A AMC**

INT CL<sup>5</sup> **G06F 12/12**

**ONLINE : WPI**

(54) User-defined priority for cache replacement mechanism.

(57) An n-way set-associative cache (where n is an integer greater than 1), includes a replacement mechanism for selecting a cache line for replacement. Each cache line has an associated priority tag indicating a user-defined priority for that cache line. The replacement mechanism comprises means for selecting a cache line with the lowest user-defined priority in a current set of cache lines, and means (e.g. based on recency of usage) for choosing between cache lines of equal priority if there is more than one cache line with the lowest user-defined priority in the set.




FIG 1

GB 2 276 962 A




FIG 1

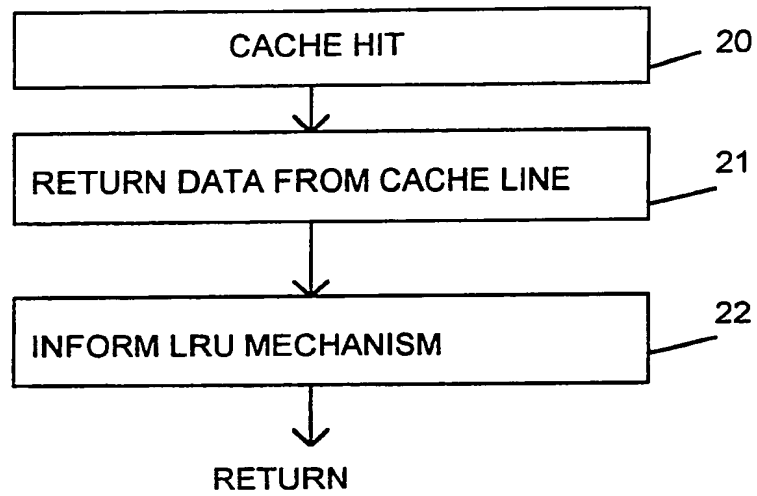


FIG 2

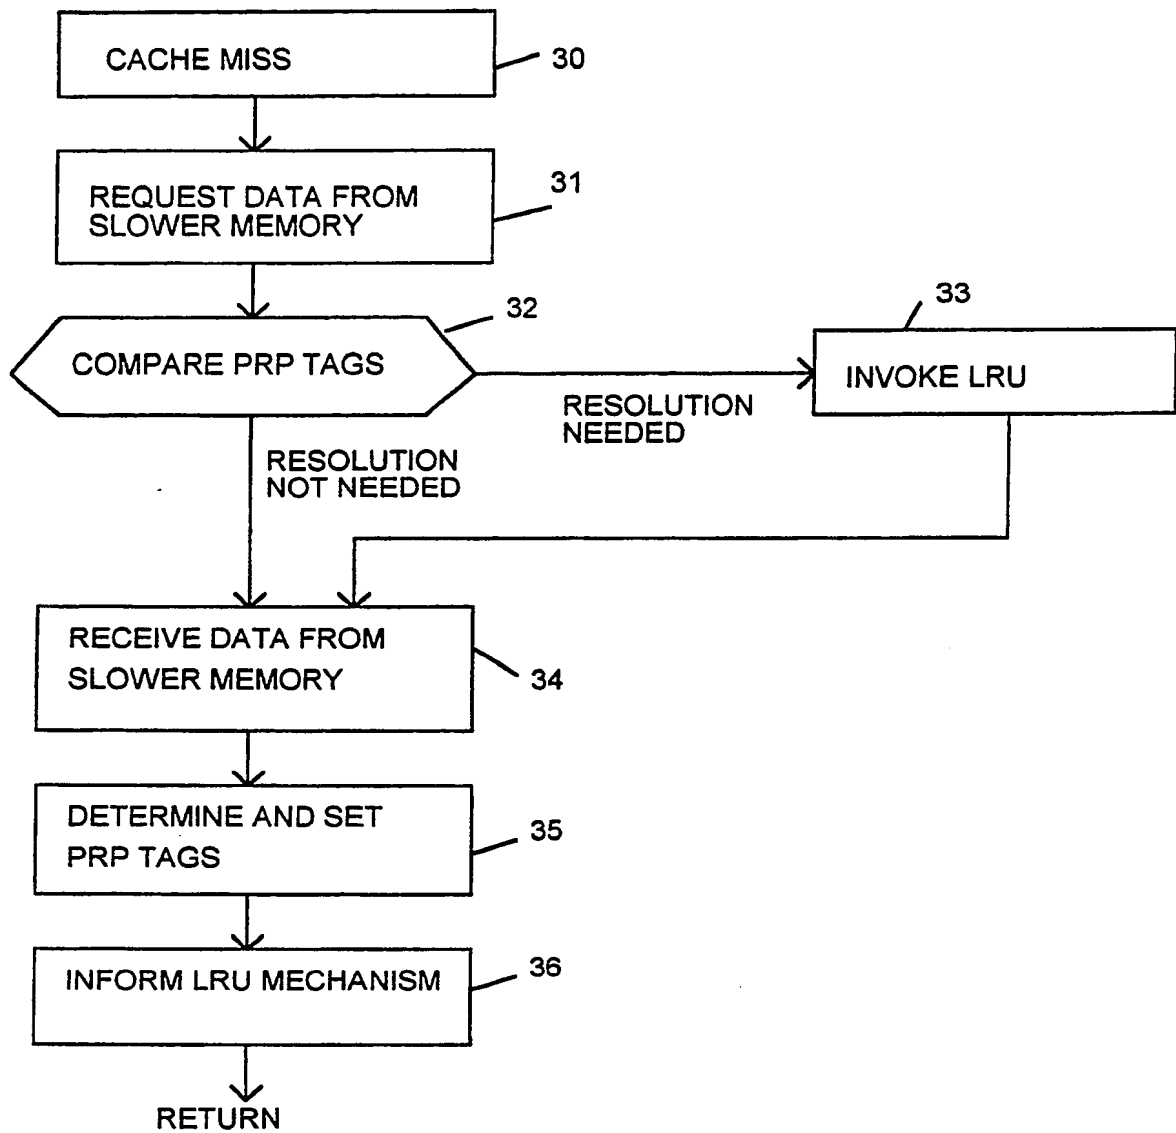


FIG 3

CACHE REPLACEMENT MECHANISMBackground to the Invention

This invention relates to set-associative cache memories.

In computer systems, it is well known to employ one or more cache memories of various sizes. The aim is to keep the most useful data in a small, fast cache in order to avoid having to retrieve the data from the larger, slower RAM. It is common to design levels of caching of different sizes and speeds.

If the required data is in a cache, it is said that a "hit" has occurred, otherwise a "miss" has occurred. The percentage of misses is called the "miss rate".

Apart from the cache size, there are two major design decisions when implementing a cache:

(1) The number of cache elements scanned simultaneously, sometimes called the "set associativity" of the cache. If just one element at a time is scanned, the cache is referred to as direct mapped. If  $n$  elements at a time are scanned (where  $n$  is greater than 1) the cache is referred to as an  $n$ -way set-associative cache. The usual choice for the value of  $n$  is 2 or 4. If all the cache is scanned simultaneously, it is referred to as fully associative. The miss rates decrease, in general, as the set associativity increases. However, the cost of implementation increases as set associativity increases.

(2) The method used to decide which cache element scanned

within the cache to replace with the desired data on a cache miss, called the cache replacement policy. This has no meaning for a direct-mapped cache since there is only one place to put the desired data. The two standard replacement methods used are: "random replacement", when the desired data is placed in one of the scanned cache elements at random, and "least recently used (LRU) replacement", when the scanned element which has been accessed least recently is replaced by the desired data. LRU replacement usually delivers the smallest miss rate but is more expensive to implement.

The object of the present invention is to provide a new cache replacement mechanism, which is potentially more efficient than these known replacement mechanisms.

#### Summary of the Invention

According to the invention there is provided an n-way set-associative cache (where n is an integer greater than 1), including a replacement mechanism for selecting a cache line for replacement, characterised in that each cache line has an associated priority tag indicating a user-defined priority for that cache line, and the replacement mechanism comprises means for selecting a cache line with the lowest user-defined priority in a current set of cache lines, and means for choosing between cache lines of equal priority if there is more than one cache line with said lowest user-defined priority in said set.

The invention thus provides a priority replacement policy (PRP), which replaces data in cache lines primarily according to user-defined priorities and secondarily according to an alternative replacement policy. The alternative replacement policy is used to choose between the data in two cache lines that have the same user-defined priority. The alternative replacement policy is said to resolve the replacement choice between the equal priority data within the currently

addressed cache lines. This alternative replacement policy may, for example, be a least-recently-used replacement policy, or a random selection.

This invention enables the data or instructions associated with a given process, application, user or user group to have relative priority within the cache. The highest priority data or instructions stay in the cache for as long as possible.

#### Brief Description of the Drawings

Figure 1 shows a cache system with priority tags.

Figure 2 is a flow chart indicating the operation of the cache system on a cache hit.

Figure 3 is a flow chart indicating the operation of the cache system on a cache miss.

#### Description of an Embodiment of the Invention

One cache system in accordance with the invention will now be described by way of example with reference to the accompanying drawings.

Referring to Figure 1, the cache system comprises a 4-way set associative cache comprising four cache data RAMs 10 and four priority tag RAMs 12. The tag RAMs contain a user-defined priority tag for each line of data in the cache. The priority may be defined explicitly, or inherited implicitly from the data's process, application, user or user's group.

The cache system also comprises a least-recently-used (LRU) replacement mechanism 14 and a priority replacement policy (PRP) mechanism 16. The LRU mechanism keeps recency of usage information relating to each cache line, and may be

conventional. The operation of the PRP mechanism will be described below.

An input memory address is received in an address register 18. This address is hashed by a hashing circuit 19 and then applied in parallel to the four cache data RAMs, so as to address one line from each RAM. The contents of the four addressed cache lines are examined to see if the desired data is resident in the cache.

Referring to Figure 2, if one of the addressed cache lines contains the desired data, then there is a hit (20) and the desired data can be immediately accessed (21) from the cache. The LRU mechanism 14 is informed (22) so that it can update the recency-of-usage information for that cache line.

Referring to Figure 3, if there is a cache miss (30) the desired data is requested (31) from slower memory. The PRP mechanism 16 then compares (32) the priority tags associated with the four addressed cache lines, to determine which of the four addressed cache lines is of lowest priority. If only one of the four addressed cache lines has this lowest priority, that line is chosen to receive the desired data from the slower memory. If on the other hand more than one data line has this lowest priority, the LRU mechanism 14 is invoked (33) to resolve the replacement choice.

When the required data is received (34) from slower memory it is written into the cache line selected for replacement. The value of the priority tag of the data is then determined (35), for example from the priority of its process, stored in a process block. This priority tag is written into the corresponding location of the priority tag RAM 12. The LRU mechanism is then informed (36) of the identity of the cache line into which the new data has been written, so that the LRU mechanism can update the recency-of-usage information for that line.



CLAIMS

1. An n-way set-associative cache (where n is an integer greater than 1), including a replacement mechanism for selecting a cache line for replacement, characterised in that each cache line has an associated priority tag indicating a user-defined priority for that cache line, and the replacement mechanism comprises means for selecting a cache line with the lowest user-defined priority in a current set of cache lines, and means for choosing between cache lines of equal priority if there is more than one cache line with said lowest user-defined priority in said set.
2. A cache according to claim 1 wherein said means for choosing between cache lines of equal priority comprises a least-recently-used replacement mechanism.
3. A cache system substantially as hereinbefore described with reference to the accompanying drawings.

**Patents Act 1977**  
**Examiner's report to the Comptroller under Section 17**  
**(The Search report)**

Application number  
GB 9404866.7

- 6 -

**Relevant Technical Fields**

- (i) UK Cl (Ed.M)      G4A AMC  
(ii) Int Cl (Ed.5)      G06F 12/12

Search Examiner  
MR S J PROBERT

Date of completion of Search  
25 MAY 1994

**Databases (see below)**

(i) UK Patent Office collections of GB, EP, WO and US patent specifications.

(ii) ONLINE: WPI

Documents considered relevant following a search in respect of Claims :-  
1-3

**Categories of documents**

- |   |   |
|---|---|
| <b>X:</b> Document indicating lack of novelty or of inventive step.   | <b>P:</b> Document published on or after the declared priority date but before the filing date of the present application.        |
| <b>Y:</b> Document indicating lack of inventive step if combined with one or more other documents of the same category. | <b>E:</b> Patent document published on or after, but with priority date earlier than, the filing date of the present application. |
| <b>A:</b> Document indicating technological background and/or state of the art.   | <b>&amp;:</b> Member of the same patent family; corresponding document.   |

Category	Identity of document and relevant passages	Relevant to claim(s)
X	EP 0391871 A2 (IBM) See pages 1-3	1-3

**Databases:** The UK Patent Office database comprises classified collections of GB, EP, WO and US patent specifications as outlined periodically in the Official Journal (Patents). The on-line databases considered for search are also listed periodically in the Official Journal (Patents).